

Report Documentation Page			Form Approved OMB No. 0704-0188					
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>								
1. REPORT DATE 28 MAY 2014	2. REPORT TYPE	3. DATES COVERED 00-00-2014 to 00-00-2014						
4. TITLE AND SUBTITLE Evaluation of long term performance of continuously running atomic fountains			5a. CONTRACT NUMBER					
			5b. GRANT NUMBER					
			5c. PROGRAM ELEMENT NUMBER					
6. AUTHOR(S)			5d. PROJECT NUMBER					
			5e. TASK NUMBER					
			5f. WORK UNIT NUMBER					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) United States Naval Observatory, Washington, DC, 20392			8. PERFORMING ORGANIZATION REPORT NUMBER					
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)					
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)					
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited								
13. SUPPLEMENTARY NOTES Metrologia 51 (2014) 263?269								
14. ABSTRACT								
15. SUBJECT TERMS								
16. SECURITY CLASSIFICATION OF: <table border="1"> <tr> <td>a. REPORT unclassified</td> <td>b. ABSTRACT unclassified</td> <td>c. THIS PAGE unclassified</td> </tr> </table>			a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 264	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified						

Evaluation of long term performance of continuously running atomic fountains

Steven Peil, James L Hanssen, Thomas B Swanson, Jennifer Taylor and Christopher R Ekstrom

United States Naval Observatory, Washington, DC 20392, USA

E-mail: steven.peil@usno.navy.mil

Received 7 March 2014, revised 22 April 2014

Accepted for publication 25 April 2014

Published 28 May 2014

Abstract

An ensemble of rubidium atomic fountain clocks has been put into operation at the US Naval Observatory (USNO). These fountains are used as continuous clocks in the manner of commercial caesium beams and hydrogen masers for the purpose of improved timing applications. Four fountains have been in operation for more than two years and are included in the ensemble used to generate the USNO master clock. Individual fountain performance is characterized by a white-frequency noise level below 2×10^{-13} and fractional-frequency stability routinely reaching the low 10^{-16} s. The highest performing pair of fountains exhibits stability consistent with each fountain integrating as white frequency noise, with Allan deviation surpassing 6×10^{-17} at 10^7 s, and with no relative drift between the fountains at the level of 7.5×10^{-19} /day. As an ensemble, the fountains generate a timescale with white-frequency noise level of 1×10^{-13} and long-term frequency stability consistent with zero drift relative to the world's primary standards at 1×10^{-18} /day. The rubidium fountains are reported to the BIPM as continuously running clocks, as opposed to secondary standards, the only cold-atom clocks so reported. Here we further characterize the performance of the individual fountains and the ensemble during the first two years in an operational environment, presenting the first look at long-term continuous behavior of fountain clocks.

Keywords: atomic clock, rubidium fountain, operational atomic fountain

(Some figures may appear in colour only in the online journal)

1. Introduction: beyond caesiums and masers

The performance of atomic clocks has improved more than seven orders of magnitude since the first caesium clock was demonstrated by Essen and Parry over 50 years ago [1, 2]. In the past two decades, improvement has been driven by several major scientific developments. Most recently, optical frequency combs, narrow line width lasers, and methods of confining atoms and ions without perturbing the clock frequency have led to optical clocks, with reference frequencies and corresponding Q s orders of magnitude higher than previous systems. Prior to optical clocks, the ground breaking techniques of laser cooling and trapping of atomic gases brought about dramatic improvements in atomic clock performance by enabling the creation of ultra-cold samples of atoms, providing longer interaction times and reduced systematic uncertainties. The first clocks based on cold atoms were atomic fountains, which are in some sense the natural

progression of atomic beam clocks. Atomic fountains are the basis of contemporary primary frequency standards, in which the microwave frequency of the hyperfine transition in caesium-133 is used as the reference for a clock [3].

As the most precise and accurate frequency standards, cold-atom clocks (optical or microwave) find application in frequency metrology [2], precise tests of fundamental symmetries [4, 5], and as test beds for quantum scattering [6], entanglement [7] and information processing [8]. However, the application of atomic frequency standards as clocks for timing applications is still dominated by clocks based on older but more mature technology, with far greater levels of engineering and evaluation. The workhorses of the timing community, commercial caesium beams and hydrogen masers, which do not incorporate any laser technology, run continuously and require little maintenance and upkeep, often running for years without requiring service. International Atomic Time (TAI) utilizes data from hundreds of these clocks,

and many timing labs rely exclusively on commercial caesiums and hydrogen masers for their local ensembles.

It is to be expected that the best optical clocks will require significant engineering before being compatible with user-free operation in a timing ensemble. (In fact, many of the best accuracy and stability evaluations carried out with optical clocks use measurements made at optical frequencies, without dividing the signal down to RF, which would be a requirement for most timing applications [9].) On the other hand, atomic fountains are being used much more regularly for timing applications. As primary frequency standards, caesium fountains have been contributing to TAI for close to a decade, with the number and precision of reports increasing over time, and recently LNE-SYRTE began reporting a rubidium fountain to the BIPM as a secondary standard [10]. And many timing labs are relying more on fountain clocks for their local timescales, with atomic fountains in use or under development in at least ten nations worldwide.

Even with this increased use, the role of fountain clocks for timing has typically been different from the commercial clocks that dominate (in number) timing ensembles. Serving as primary frequency standards, the fountains provide calibrations of the frequencies of other clocks. In this role, they do not necessarily need to run continuously for months or years at a time without user intervention, though some have achieved close to continuous operation [11, 12]. And for a local timescale, it can suffice to use only a single primary standard, which ensures a constant long-term frequency, at least on the time frame of the accuracy evaluations. For example, the PTB has produced an excellent timescale based on a single caesium fountain [11].

Still, there are advantages to having multiple clocks. More clocks translate to more robustness and reliability. And it is sometimes the case that a fountain clock is characterized by comparing to a second one. For these reasons, most institutions with fountains have or are developing at least two; LNE-SYRTE has three and NPL is working on a third (in each case one of the three fountains uses rubidium). Additional advantages of multiple clocks, particularly for timekeeping, are related to reducing noise. The white-frequency-noise level of an ensemble with N clocks decreases as \sqrt{N} . And since continuously running atomic clocks exhibit non-stationary behavior, timing ensembles based on such clocks require multiple devices to optimize performance.

The US Naval Observatory (USNO), the largest contributor of clock data to the BIPM, one of the official sources of time for the United States and the official time and frequency reference for critical infrastructure such as the Global Positioning System, has designed and built an ensemble of atomic fountain clocks that has been in operation for more than 2.5 years. The ensemble consists of four rubidium fountains housed in Washington, DC¹. These clocks run continuously, with little operator intervention required, and are used to enhance the local timing ensemble, contributing to the USNO master clock in a manner similar to commercial caesium beams and hydrogen masers. The specific role of the fountains is to provide a superior long-term frequency

reference, in the past solely the job of the commercial caesiums, while the short-term reference continues to be provided by an ensemble of hydrogen masers.

All four rubidium fountains have been regularly reported to the BIPM for over two years as continuously running clocks, the first cold-atom clocks so reported. Here we present the performance of these clocks over the first two years of operation, the first characterization of continuously running atomic fountains over such time frames.

2. Overview of fountain operation

The four fountains in operation at USNO are designated NRF2, NRF3, NRF4 and NRF5. The systems were built in two generations, according to slightly different designs [13–15]. All four fountains began continuous operation in a dedicated clock facility in March 2011. Since then, some of the improvements in design used for the second generation, NRF4 and NRF5, have been incorporated as retrofits to the first generation, NRF2 and NRF3. Additionally, both routine and unexpected maintenance were required on occasion. Since December 2011, fountain data for all four devices have been reported without interruption to the BIPM, as clock type 93. In May 2012, after the traditional evaluation period, the fountains were weighted in TAI for the first time, and all four clocks have received maximum weight in every evaluation since.

Each fountain has a local oscillator (LO) comprising a quartz crystal phase-locked to a dedicated maser, which serves as the reference for the 6.8 GHz drive for rubidium spectroscopy as well as the reference for a precise frequency synthesizer, an auxiliary output generator (AOG). Every 20 s, the average relative frequency of the fountain and maser is used to steer the AOG output, providing a continuous nominal 5 MHz signal that reflects the fountain frequency and is measured against the USNO master clock and other clocks in the ensemble.

During each fountain cycle, auxiliary data such as the number of atoms participating in the frequency measurement and the fraction of atoms making the transition between clock states are used to determine whether the fountain is operating normally. If these data fall outside of a predetermined healthy range, the fountain automatically switches to holdover mode, in which the maser serves as a flywheel; in this case, the control computer applies a steer that is the median frequency difference between the fountain and maser over the previous hour. The same holdover steer is applied until the auxiliary data reflect healthy values and normal operation resumes.

The robustness of the fountains and the ability to rely on a maser as a flywheel have resulted in a high uptime over the two year period. The percentage of time that each fountain has generated a good, steered output is 99.2%, 99.7%, 98.5% and 100% for NRF2 through NRF5. A ‘good output’ can include brief intervals where the system is in holdover; so long as the interval corresponds to an averaging time for which the relative frequency variations of the maser and fountain are still white (i.e. the frequency fluctuations are still limited by fountain performance). Periods in holdover long enough that the steered output reflects maser behavior are treated as ‘down

¹ Two additional fountains have been built and tested and are in preparation for use at USNO’s Alternate Master Clock in Colorado.

time' and this is not included in analysis of fountain behavior; it is handled in the data analysis by interpolating the frequency record and integrating to generate the phase record.

The quickest degradation of performance of the (nominal) 5 MHz output signal is the occurrence of a problem with the processor controlling fountain operation; if the steering algorithm is interrupted, the AOG will no longer be updated, and its output frequency will reflect a steer determined from 20 s of averaging (rather than the holdover value). Processor refresh is usually a simple process. Problems that are more difficult to solve, but do not show up as quickly in the steered output, include shutter failure and diode laser aging. Aside from these three problems that have arisen, routine maintenance is required on occasion. This consists almost exclusively of occasional tweaks to optical alignment, sometimes not needed for months at a time. Very rarely an adjustment to the voltage-controlled oscillator on the reference maser or AOG is required to maintain phase lock. There is no maintenance required for the laser frequency locks, which perform without incident.

3. Clock performance

3.1. White frequency noise

Perhaps in part due to the refinement in design, the highest performing fountains have been NRF4 and NRF5. Each device exhibits a white-frequency noise level of 1.8×10^{-13} and excellent long-term stability. (There is an obvious enhancement to the microwave chain which should improve the short-term stability to a level of about $1.4 \times 10^{-13}/\sqrt{\tau}$, which was demonstrated in an engineering prototype, NRF1.)

In figure 1(a), the relative phase of NRF4 and NRF5 is plotted over an uninterrupted period of more than 1 year, after removing a relative frequency of 4.3×10^{-16} and a meaningless relative phase. Residual peak-to-peak phase deviations stay within about 1 ns for the entire interval. A limit on the relative frequency drift can be obtained by fitting a line to the frequency, with the result that the relative drift is zero at the level of $7.5 \times 10^{-19}/\text{day}$ ($2.7 \times 10^{-16}/\text{year}$).

A sigma-tau plot, shown in (b), includes overlapping Allan deviation, total deviation and Theo statistics, along with a white-frequency noise reference line. There is no indication of deviation from white-frequency noise behavior out to the total averaging time of more than 8 months. The average (in)stability for each fountain is consistent with an overlapping Allan deviation of 6.8×10^{-17} at 7×10^6 s, a total deviation of 4.6×10^{-17} at 1.5×10^7 s and a Theo of 3.8×10^{-17} at 2.2×10^7 s (values correspond to the white-frequency reference level divided by $\sqrt{2}$, not to the data points which lie below this line). Figure 1(c) shows the time deviation (TDEV) between the two fountains for the ~ 1 year interval; the average TDEV for each fountain is about 20 ps at a day.

It is likely that this represents record performance for microwave clocks.

3.2. Non-stationary behavior

The data analyzed in figure 1 correspond to an interval when both fountains exhibited excellent stability, which also

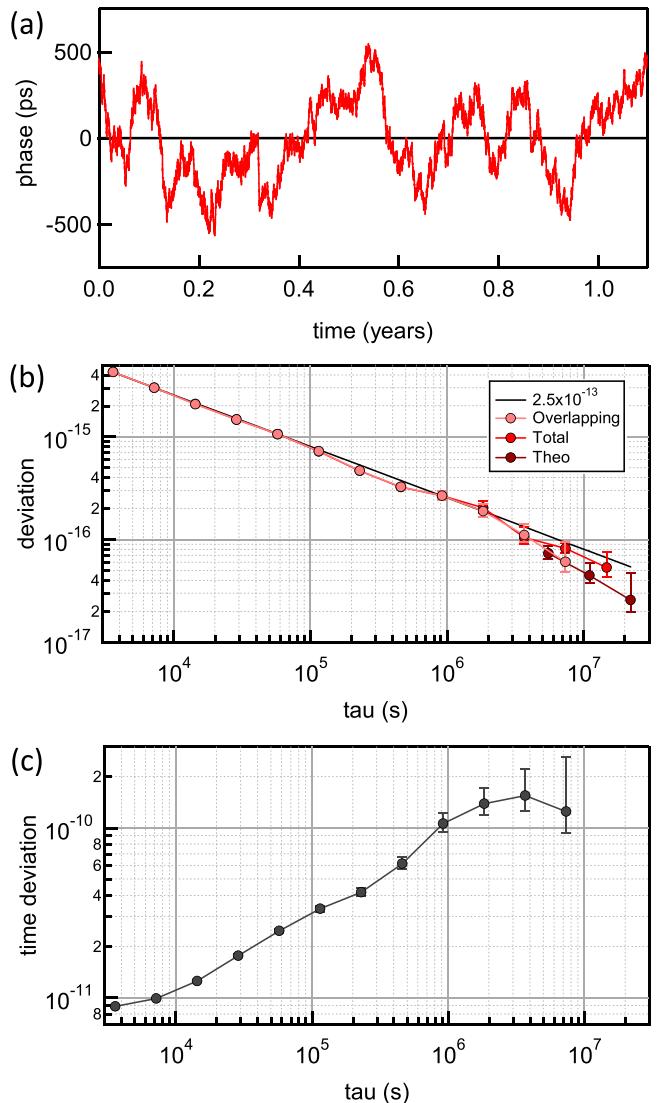


Figure 1. Performance of NRF4 versus NRF5 for 1.1 year interval. (a) Residual difference in phase after removing a single frequency difference. (b) Sigma-tau plot for data in (a). Plot includes overlapping Allan deviation, total deviation, and Theo statistics. The white-frequency noise reference line corresponds to an average level of 1.8×10^{-13} for each fountain. (c) Time deviation for data in (a).

coincided with a period of time during which there was little or no user intervention. There were times during the two years when some final refinements were made to NRF4, upgrades were made to NRF2 and NRF3, and maintenance consisting of laser and shutter replacement were carried out on NRF2 and NRF3. When the entire two-year data set is analyzed without any restrictions, and with no corrections or re-evaluations of the frequency across the service or upgrade, we see long-term behavior that deviates from white frequency noise.

Figure 2 shows the overlapping Allan deviations for all pair-wise comparisons using the entire two-year data set. The long-term behavior in all cases shows similar non-stationary behavior. When the NRF4 versus NRF5 comparison in (a) is extended to longer averaging times using the Theo deviation (dark symbols), the most obvious fit is to random-walk frequency noise, with a $\sqrt{\tau}$ dependence. The reference line in

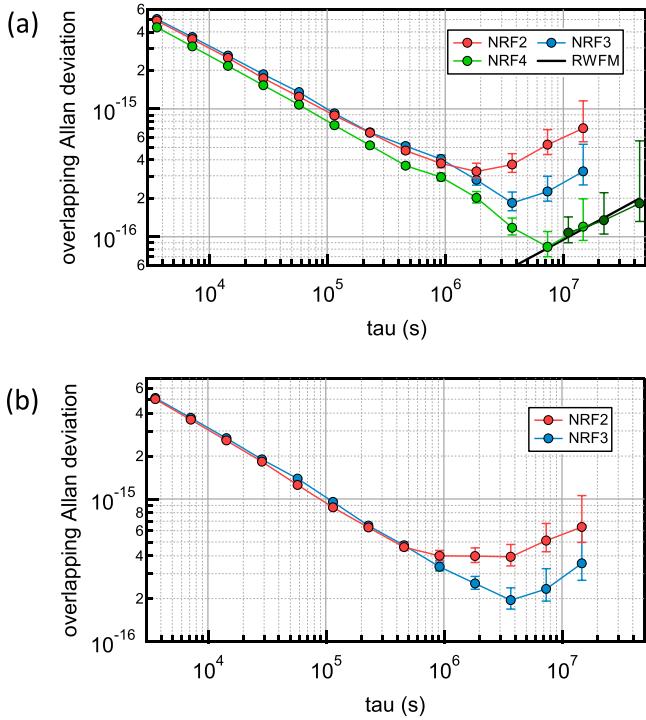


Figure 2. Summary of pair-wise fountain stabilities over first two years of operation. (a) Stabilities of NRF2, NRF3 and NRF4 measured against NRF5. RWFM—random walk frequency modulation. (b) Stabilities of NRF2 and NRF3 measured against NRF4.

the figure corresponds to a random-walk level of $2 \times 10^{-20} \sqrt{\tau}$ per fountain. But, as we now show, the behavior leading to these stability plots is entirely consistent with discrete changes in frequency, at a rate that makes it impossible at this point to distinguish between inherent random-walk frequency noise and frequency changes arising from user intervention.

In figure 3(a), we show the relative phase of NRF4 and NRF5 over the entire two year period. The record of the relative phase shows two different linear regions with different slopes, indicating that the relative frequency changed (the data in figure 1 correspond to the second frequency shown). This bears out when we adjust the data starting from 0.9 years (MJD 55975) to the end by removing a single frequency, of size $< 3 \times 10^{-16}$, in order to make the last part of the frequency record match the first. The subsequent data look white, as illustrated by the sigma-tau plot in figure 3(b), indicating that the non-stationary behavior is indeed consistent with a single frequency change.

Similar behavior was observed for NRF3. In figure 4(a), we show the relative phase of NRF3 and NRF5 over the entire two year period. Here, the relative frequency appears to take on three different values. Adjusting the data at the two times when the frequency changed (0.6 years (MJD 55840) and 1.1 years (MJD 56030)), thereby adding two frequency adjustments, the resulting behavior is consistent with integration as white-frequency noise, as shown by the sigma-tau plot in figure 4(b). So for NRF3, NRF4 and NRF5, the long-term behavior under continuous operation is consistent with rare, discrete frequency changes, at an average rate of one every two years and an average size of 5.4×10^{-16} .

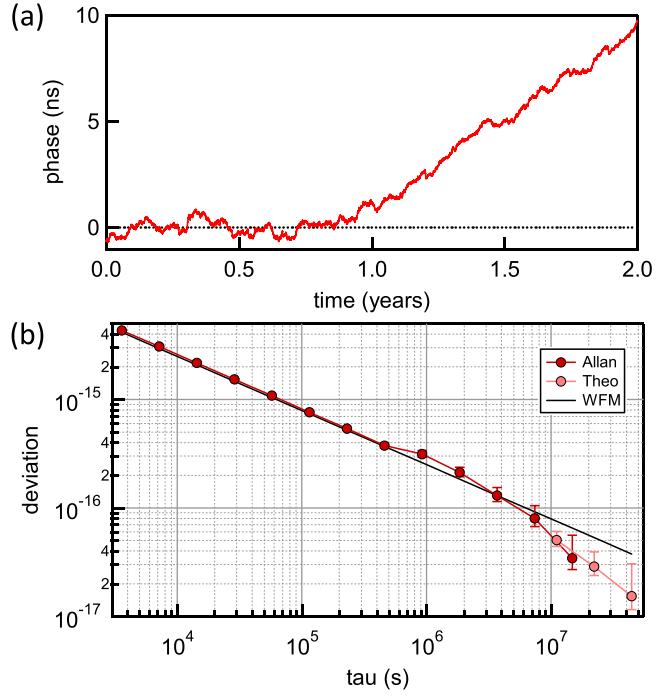


Figure 3. (a) Phase record for NRF4 versus NRF5, over the first two years of operation (from MJD 55637 to 56368). (b) Overlapping Allan deviation and Theo for data set created when a single frequency is removed from the latter part of the phase record, whitening the data. WFM—white frequency modulation.

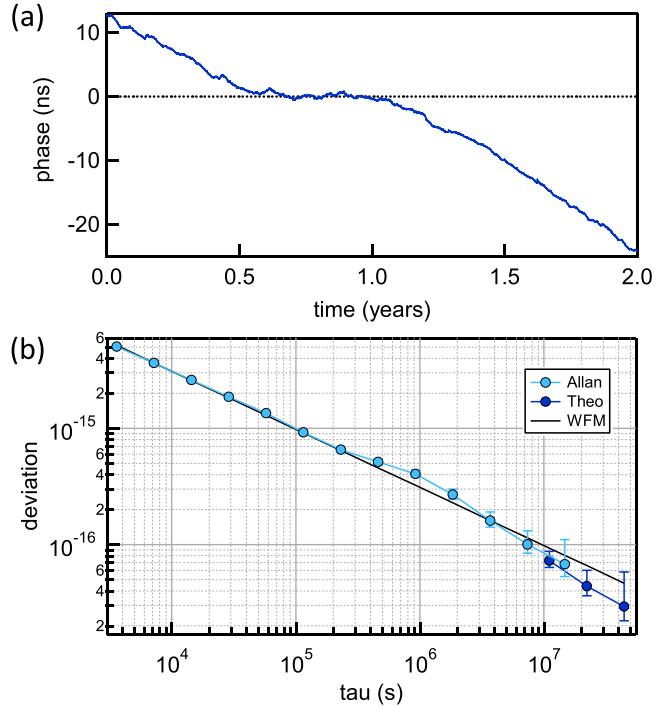


Figure 4. (a) Phase record for NRF3 versus NRF5. (b) Overlapping Allan deviation and Theo for data set created when frequency adjustments are made to the data at two points.

We can unambiguously assign the frequency changes in figure 4 to NRF3, as they appear as well in comparisons with NRF4. The frequency change in figure 3 is more difficult to assign to one fountain or the other, because the other two

fountains do not support a particular position. We suspect that the frequency change can be assigned to NRF4, however, because the time at which the change occurred coincides within the error bar to a time at which some finishing touches were completed on NRF4. Similarly, one of the changes in frequency exhibited in figure 4(a) occurred close in time to one of the upgrades made to bring the design of NRF3 closer to that of the second generation. It is impossible to conclude that these frequency changes were caused by user intervention, but it is a possibility, and it is something we intend to continue analyzing.

Almost all of the most likely sources of significant frequency shifts in a fountain have been ruled out as drivers of the observed behavior—the magnetic field is regularly measured and shows no correlated changes; light shifts do not seem to be a problem; the ambient temperature is monitored and is constant. It is hard to conclusively rule out a microwave effect, particularly within the microwave chain. It would be interesting to compare to primary standards, in terms of whether biases measured in accuracy evaluations change occasionally over time.

Because these frequency changes are rare, it is straightforward to recharacterize a fountain using the others as a reference. This introduces the possibility of producing a paper fountain timescale that conceivably integrates as white-frequency noise to an order of 1×10^{-17} . We plan to consider this further.

The fourth fountain, NRF2, has exhibited the least stable behavior thus far. It has shown discrete frequency changes, of larger size and at a higher rate than those discussed above, as well as intervals that more likely corresponded to drift. Some of the problematic behavior in NRF2 does not correlate with episodes of significant intervention. We hope to be able to improve the performance of NRF2 and in the process identify the cause of at least some of this non-stationary behavior.

Finally, we point out that these changes in the output frequencies of the fountains are not observable using other clocks or timescales at USNO. Neither the caesium ensemble, nor the maser ensemble, nor individual masers independently show that a fountain has changed frequency; these events are only observed with inter-fountain comparisons.

4. Role in local timescale, UTC(USNO)

For years the USNO master clock, the physical representation of UTC(USNO), has been generated by incorporating the long-term frequency stability of commercial caesium beams with the superior short-term stability of hydrogen masers. Maser are well known to display time-variable drift and other behavior different from white-frequency noise after integrating to a level between low 10^{-15} s and high 10^{-16} s. Each maser is re-characterized using the caesiums; the sooner that non-stationary behavior is detected, the better. One role of the rubidium fountains is to improve on the process of maser re-characterization.

In figure 5 we show the Allan deviation of an exceptionally stable maser versus a crude fountain timescale, created by averaging the frequencies of the four rubidium fountains. The

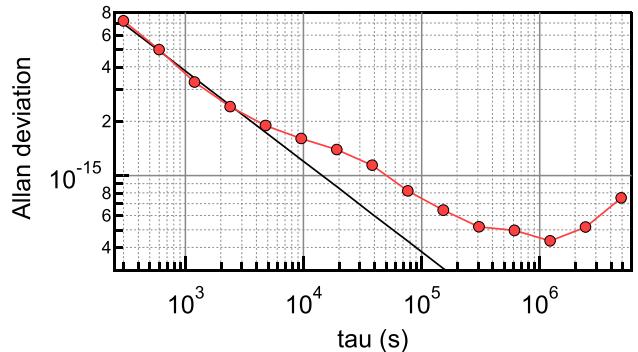


Figure 5. Overlapping Allan deviation of the average of the four fountains versus a hydrogen maser (circles). The white-frequency noise reference line (solid line) is at 1.2×10^{-13} , corresponding to 1×10^{-13} for the fountain timescale and 6×10^{-14} for the maser.

short term relative stability of 1.2×10^{-13} at 1 s is consistent with a white frequency noise level of 6×10^{-14} for the maser and 1×10^{-13} for the fountain timescale, as expected for four fountains with an average white-frequency noise of 2×10^{-13} . The white frequency noise level achieved with the USNO ensemble of about 70 commercial caesium beams is about 1×10^{-12} . So the fountain ensemble shows a factor of 10 improvement in short-term stability, which, in terms of averaging time, translates to a factor of 100 reduction. Even a single fountain is more than 50 times better than the entire ensemble of caesiums.

As stated previously, it should be possible to improve the short-term stability of each fountain to about 1.4×10^{-13} , making the fountain timescale white frequency noise level 7×10^{-14} . We expect that this would be the limit for our four-fountain ensemble. We have investigated using the same LO for a pair of fountains to enable reduction of the LO noise contribution. Removing the entire LO noise contribution by differencing the fountain frequencies is possible for a paper clock and is useful for measuring atomic physics effects more quickly, but it cannot be incorporated into a physical output. To generate a physical output with reduced LO noise, we implemented what we call ‘cooperative steering’, where two fountains steer the same LO. By synchronizing the load-measurement cycles so that the LO is always being measured and compensated for, the additional noise from not continually sampling the LO, the Dick effect, can be reduced [16]. Because of the high quality of the LOs used for the rubidium fountains, the reduction in the Dick effect does not improve the stability greatly, while using a single LO for different fountains increases risk significantly. This increase in operational risk led us to abandon this technique.

As confidence in the reliability of the fountains has been gained, the clocks are now being used operationally as the long-term reference in the generation of UTC(USNO). The specific algorithm for including the fountains as well as a quantitative assessment of their impact on UTC(USNO) is an area of ongoing research [17].

5. Stability versus primaries contributing to TAI

Eight caesium fountains worldwide contributed on some basis to TAI during the two years we are considering.

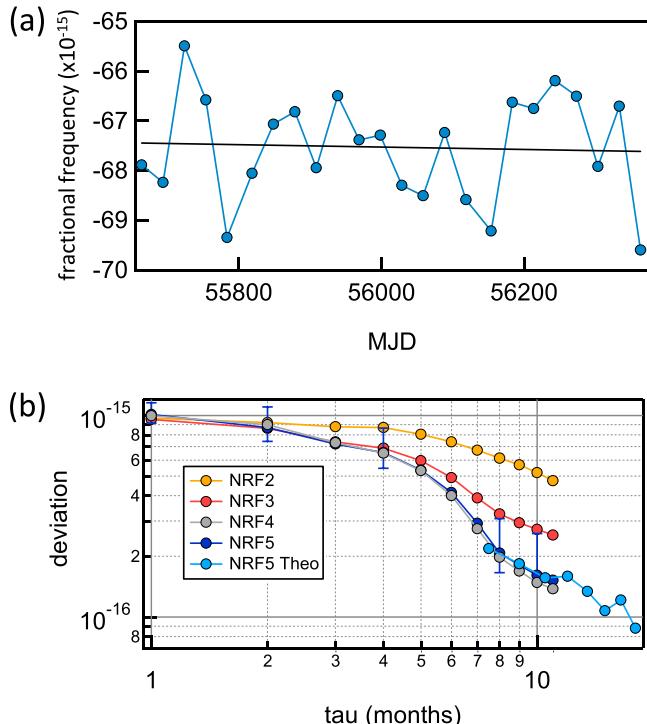


Figure 6. (a) Relative frequency of NRF5 and the average of the caesium fountains contributing to TAI for a particular interval. A line fit to the data gives a slope of $0.2 \pm 1.0 \times 10^{-18}/\text{per day}$, demonstrating that NRF5 exhibits zero drift with respect to the world's primaries at the level of $1.2 \times 10^{-18}/\text{per day}$. (b) Total deviation for each individual fountain versus the caesium fountains. Theo deviation and sample error bars are included for NRF5.

Perhaps the best metric of long-term frequency stability of the USNO rubidium fountains is by comparing to these primary standards². We do this by using the USNO master clock, the physical realization of UTC(USNO), as a transfer oscillator to allow us to determine the frequency difference between a fountain and TAI, $\nu_{\text{Rb}} - \nu_{\text{TAI}}$. Along with the values of the average frequency of the primaries compared to EAL, $\nu_{\text{Cs}} - \nu_{\text{EAL}}$, and the steers of EAL to TAI, $\nu_{\text{EAL}} - \nu_{\text{TAI}}$, which are published monthly by the BIPM in their *Circular T* report, the frequency of a fountain or fountain timescale compared to the primaries can be obtained. Figure 6(a) shows the monthly average of the frequency of NRF5 versus the average of the primaries. The line in the plot is a linear fit, $0.2 \pm 1.0 \times 10^{-18}/\text{day}$, showing that we can place a limit on the drift of NRF5 with respect to the primaries at the level of $1.2 \times 10^{-18}/\text{per day}$. In figure 6(b), a sigma-tau plot shows the stability of each fountain versus the primaries.

Because of the continuous operation of our rubidium fountains, in addition to comparing to the monthly average of the primaries, we can obtain a measurement against each individual primary standard reported or against standards at a particular institution. This allows us to analyze trends with respect to specific standards labs. This is illustrated in figure 7, where the relative frequency of NRF5 and each primary standard reported to TAI over the two years is plotted.

² The primaries data set also includes CS1 and CS2 at PTB, which are beam standards.

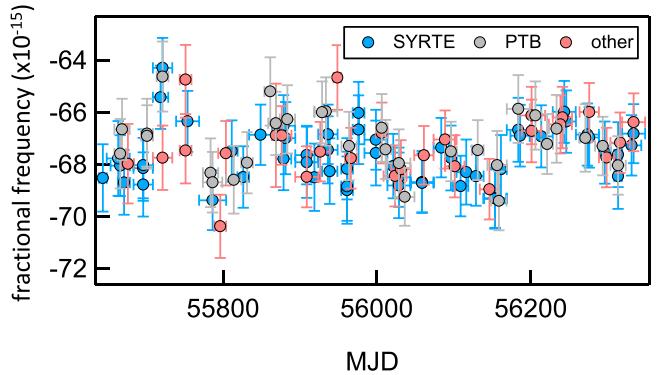


Figure 7. Relative frequency of NRF5 and each caesium fountain measurement reported to TAI.

The (three) fountains at LNE-SYRTE and (the two) at PTB are differentiated from the other labs reporting (NIST, NPL and one report from NICT).

6. Conclusion

To summarize, the USNO has put into operation an ensemble of continuously running rubidium fountain clocks to incorporate into timing applications and timescale generation in a manner similar to hydrogen masers and commercial caesium beam clocks. We have analyzed performance of the four fountains in the ensemble over the first two years of operation, demonstrating intervals reflecting record performance for microwave clocks and observing rare, discrete frequency changes. Comparison to the world's primary standards suggest that the rubidium fountains provide a stable, long-term frequency reference on par with any local reference.

Acknowledgments

We benefited from assistance from many members of the USNO Time Service Department, particularly Paul Koppang, Jim Skinner and Demetrios Matsakis. Scott Crane played a significant role in the design and construction of the fountains. Atomic fountain development at USNO has been funded by ONR and SPAWAR.

References

- [1] Essen L and Parry J V L 1955 *Nature* **176** 4476
- [2] Rosenband T et al 2008 *Science* **319** 1808
- [3] Wynands R and Weyers S 2005 *Metrologia* **42** S64
- [4] Peil S, Crane S, Hanssen J L, Swanson T B and Ekstrom C R 2013 *Phys. Rev. A* **87** 010102(R)
- [5] Wolf P, Chapelet F, Bize S and Clairon A 2006 *Phys. Rev. Lett.* **96** 060801
- [6] Legere R and Gibble K 1998 *Phys. Rev. Lett.* **81** 5780
- [7] Louchet-Chauvet A, Appel J, Renema J J, Oblak D, Kjaergaard N and Polzik E S 2010 *New J. Phys.* **12** 065032
- [8] Rosenband T et al 2007 *Phys. Rev. Lett.* **98** 220801
- [9] Hinkley N et al 2013 *Science* **341** 1215
- [10] Guéna J, Abgrall M, Clairon A and Bize S 2014 *Metrologia* **51** 108

- [11] Bauch A, Weyers S, Piester D, Staliuniene E and Yang W 2012 *Metrologia* **49** 180
- [12] Guéna J et al 2012 *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* **59** 391
- [13] Peil S, Crane S, Swanson T B and Ekstrom C R 2005 *Proc. 2005 IEEE Int. Frequency Control Symp. (Vancouver)* p 304
- [14] Peil S, Crane S, Swanson T and Ekstrom C R 2006 *Proc. 20th European Frequency and Time Forum (Braunschweig, Germany, 2006)* p 194
- [15] Ekstrom C, Peil S, Swanson T and Crane S 2009 *Frequency Standards and Metrology, Proc. 7th Symp. (Pacific Grove, CA)* (Singapore: World Scientific) p 308
- [16] Dick G J, Prestage J D, Greenhall C A and Maleki L 1987 *Proc. 19th Precise Time and Time Interval (PTTI) Applications and Planning Meeting (Redondo, CA)* (Washington, DC: National Aeronautics and Space Administration) p 133
- [17] Koppang P, private communication